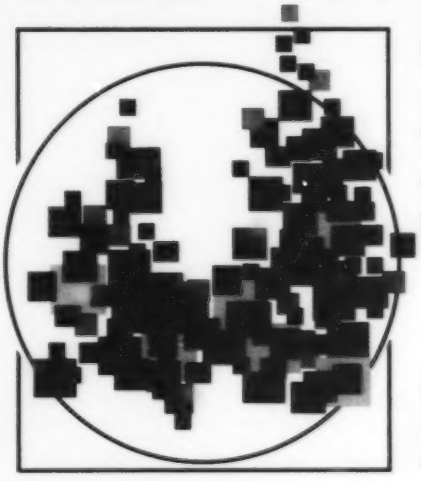

CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Calculating Surface Reflectance using a Single-Bounce Model of Mutual Reflection

Mark S. Drew and Brian V. Funt

CSS/LCCR TR90-17

July 1990

92-05798/

Calculating Surface Reflectance using a Single-Bounce Model of Mutual Reflection

Mark S. Drew
Brian V. Funt

School of Computing Science
Simon Fraser University
Vancouver, British Columbia, CANADA V5A 1S6
(604) 291-4682

Keywords: Computer vision; color perception; color constancy; finite-dimensional linear model; mutual reflection.

To be presented at ICCV '90, Osaka, Dec.4-7, 1990.

Copyright © 1990 Mark S. Drew and Brian V. Funt

Abstract

Light reflected from one surface onto a second surface changes both the intensity and spectral power distribution of light leaving the second surface. Similarly, light from the second surface illuminates the first. This mutual reflection effect can be exploited by examining pixels where interreflection is and is not present. From these measurements several intrinsic properties can be determined: the reflectance of each surface, the spectral power distribution of the incident illumination, and some constraints on the physical configuration of the two surfaces. We use finite dimensional linear models for the ambient illumination and for surface spectral reflectance, with m basis functions for illumination and n for surfaces. Then examining p sensor values (e.g. RGB values) we find that if p satisfies the condition $p \geq (2n + m) / 3$ we can solve for finite dimensional model descriptors of both surfaces and of the ambient illumination, as well as for a form-factor stemming from the surface configuration. With n=m=3, p can also be 3. A single-bounce model of mutual reflection accounts for the most important contribution to light intensity in an interreflecting geometry. We have previously tested a one-bounce model using simulations based on input values derived from an infinite-bounce, two-zone model. The results were excellent. Here, we test the one-bounce model with data corresponding to an infinite-bounce, multi-zone simulation. The results show that surface reflectance is recovered with a reasonable degree of accuracy.

1. Introduction

Two illuminated surfaces of different reflectance may have their colors bleed into one another in regions where light reflected from one surface falls onto the other surface. This is a mutual reflection effect. E.g., in Fig.1 we show two surfaces that form an edge and the resulting interreflection effect. In this paper we examine how the lights reflected separately from each surface combine to form RGB values in the mutual reflection region. Using vector models for the illumination spectral power distribution and for surface reflectance functions we show that by looking at the extra information coming from measurements of the interreflection RGB values, in addition to sensor measurements from each surface separately, it is possible to recover the surface reflectance properties of both surfaces. Thus, we obtain a color constancy algorithm.

Maloney and Wandell [1, 2, 3] use finite dimensional linear models to recover reflectance descriptors from RGB camera responses for an image. However, several special conditions have to be true for reflectance to be recoverable from the color signal entering the camera. E.g., the number of sensor classes (e.g., 3 for RGB sensors) must be at least as great as the number of basis functions modeling the surface spectral reflectance plus 1. This means that for a (minimum) dimensionality for modeling reflectance of 3 one must somehow develop a fourth "sensor class" to provide enough information to allow a solution that disentangles illumination and surface spectral reflectance. Alternatively, one can continue using the usual 3 sensor classes, but at the expense of having only a dimensionality of 2 with which to model surface reflectance.

Here we examine sensor values ρ_k at a 3-dimensional edge which include a mutual reflection contribution. We show that the resulting bleeding of colors at the 2-dimensional image edge can be used to circumvent the "fourth sensor" problem.

Suppose we examine sensor values ρ_k, k=1..p, and try to determine the finite dimensional linear model weights, e_i, i=1..m for illumination and σ_j, j=1..n for surfaces. We can take measurements of ρ_k values from pixels where mutual reflection effects are present and compare them to values where such effects must be absent. By using such

measurements we capture sufficient information to determine ε_i and σ_j separately on each side of the edge, and hence determine intrinsic reflectance descriptors for each color patch. These measurements yield, then, a color constancy algorithm. Our algorithm finds constant color descriptors which might be useful in removing the mutual reflection component from the image. Such an image filter could be useful for shape from shading, etc. [4, 5].

The only constraint on the formulation turns out to be that the number p of sensor classes must be greater than or equal to $(2n + m)/3$, where m and n are the number of basis functions required to model illumination and surface spectral reflectance respectively. This constraint is satisfied by values $p = m = n = 3$ and hence we can keep to only 3 sensor classes while still allowing $n = 3$ for describing reflectances. These values of p, m, and n are shown to recover both surface spectral reflectance and illumination spectral curves well within acceptable errors.

This paper addresses only a small part of a large problem. We assume a preliminary edge-finding and segmentation algorithm that would identify those edges where mutual reflection might be present. Given such identification of appropriate sites our analysis concentrates on extracting any available reflectance and configuration information from mutual reflection effects.

In [6] we considered input signals made up from an infinite-bounce interrflecting edge, where it was assumed that each side can effectively be divided into two zones-- one with and one without mutual reflection effects being present. We showed there that the results of a one-bounce, two-zone model are reasonably good in most cases, with accuracies of the resulting reflectance spectra not degraded much from the accuracy of the underlying finite dimensional model. Here we test the model further by producing a simulated edge, taking into account infinite bounces again but sub-dividing each side into multiple zones, as in the Radiosity Method [7] in computer graphics. We show in Section 5 that the results are good-- we can recover reflectances on each side of the edge reasonably well from mutual-reflection-induced color gradations.

2. Finite-dimensional Models Applied to Mutual Reflection

We assume that illumination and surface spectral reflectance are each modeled to an acceptable degree of accuracy by finite dimensional linear models. Depending on the dimensionality of the set of basis functions, not all illumination spectra and surface reflectances will be well-modeled. However, we shall show below, using simulations, that for natural occurring spectra the accuracy of the solution for color constancy from mutual reflection measurements is quite good.

For reflectance basis functions S_j we use those determined by Cohen via a statistical principal component analysis of 150 Munsell chips randomly selected from a total of 433 chips [8; see also 9]. For illumination, a similar analysis was carried out by Judd et al. [10] for 622 typical daylight samples, with resulting basis functions E_i. We demonstrate below that it is possible to use $n = m = 3$ basis functions for surfaces and illumination and still have only $p = 3$ sensor classes.

Most naturally occurring daylights and many surface reflectances can be approximated reasonably well using these basis functions, so that by determining a good approximation of the weights ε_i and σ_j, we are in fact deriving an approximation of the entire illumination or reflectance spectrum. Hence what we require from a solution is actually more stringent than simply *trichromatic* color constancy, in that we seek to recover more information than just chromaticity values.

Consider the situation depicted in Fig.1. Uniform diffuse ambient illumination $E(\lambda)$, impinges on two surfaces, 1 and 2. Assume surface 1 has spectral reflectance function $S^{(1)}(\lambda)$ and surface 2 has reflectance $S^{(2)}(\lambda)$, all surfaces being ideal diffuse (i.e., Lambertian) reflectors. Then depending on the geometry, interreflection may be present as shown in Fig.1. In fact, the two surfaces need not meet at an edge for a mutual reflection effect to come into play; and we do not require the surfaces to be flat. We show ray geometry in Fig.1 to clarify the interreflection process. However, ray geometry is actually unnecessary since perfect diffuse reflectors obliterate directional information: reflections are really in all directions into the hemisphere above each surface. We assume here that the incident illumination not attributable to mutual reflection is uniform across both surfaces.

The color signal $I(\lambda)$ is the light reflected by the surface, and is given by the product of E and S. On the surface 1 side, at a point relatively far from the edge, the color signal is simply

$$I^{(1)}(\lambda) = E(\lambda)\, S^{(1)}(\lambda).$$

Similarly, on the surface 2 side we have

$$I^{(2)}(\lambda) = E(\lambda)\, S^{(2)}(\lambda).$$

Substituting the basis function decomposition, we have

$$I^{(1)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(1)}_j E_i(\lambda) S_j(\lambda)$$

$$I^{(2)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(2)}_j E_i(\lambda) S_j(\lambda)$$

using weights ε_i for illumination, and $\sigma^{(1)}_j$ and $\sigma^{(2)}_j$ for reflectance on each side of the edge.

Now consider the mutual reflection region in which part of $I^{(1)}(\lambda)$, the light reflected from surface 1, impinges onto surface 2. Similarly, part of $I^{(2)}(\lambda)$ impinges onto surface 1. The amount of $I^{(1)}(\lambda)$ that is intercepted by surface 2 is dependent on both the configuration of the surfaces and on the shape of each surface, since not every reflected ray is intercepted.

For clarity, in Fig.1 only a single interreflection is shown, but in general we consider the fraction α_{12} of $I(\lambda)$ from 1 that strikes 2, as well as the similarly defined α_{21}. Denote by $I_m^{(2)}$ the color signal from a spot on surface 2 in the mutual reflection region, so that it includes a contribution from light reflected by surface 1. It consists of two parts. The first part is $I^{(2)}(\lambda)$, the reflection of $E(\lambda)$ from surface 2. This we assume is the same as from a spot relatively far from the edge, outside the mutual reflection region. The second part is due to the light reflected from surface 1. Since the latter is presumed to come from a spot on surface 1 that is near the edge, the second part is a contribution from $I_m^{(1)}$ (cf. [7]). Hence the color signals coming from each side of the edge in the mutual reflection region are*

*We include no dependency of the surface spectral reflectance function shape on angle of incidence or angle of line of sight, although it could change by an overall multiplicative constant.

$$I_m^{(1)}(\lambda) = I^{(1)}(\lambda) + \alpha_{21} I_m^{(2)}(\lambda) S^{(1)}(\lambda)$$

$$I_m^{(2)}(\lambda) = I^{(2)}(\lambda) + \alpha_{12} I_m^{(1)}(\lambda) S^{(2)}(\lambda) \quad (1)$$

where α_{12} is not assumed to be the same as α_{21}.

In [6] we solve the above set of simultaneous equations and show that to first order they reduce to

$$I_m^{(1)}(\lambda) = I^{(1)}(\lambda) + \alpha_{21} I^{(2)}(\lambda) S^{(1)}(\lambda)$$

$$I_m^{(2)}(\lambda) = I^{(2)}(\lambda) + \alpha_{12} I^{(1)}(\lambda) S^{(2)}(\lambda) \quad (2)$$

The first order approximation simply replaces $I_m(\lambda)$ with $I(\lambda)$ on the right hand sides, thereby supressing the effects of multiple reflections. These equations, therefore, constitute a one-bounce, two-zone model of mutual reflection.

To reduce the above equations to ones involving *sensor values*, we multiply by the sensor sensitivity functions $R_k(\lambda)$ for $k = 1 .. p$ sensor classes and integrate over wavelength (e.g., p=3 for RGB signals). Then following [2] and defining the pre-calculated tensor

$$g_{ijk} = \int E_i(\lambda) S_j(\lambda) R_k(\lambda) d\lambda ,$$

the equations for $I^{(1)}(\lambda)$ and $I^{(2)}(\lambda)$ become the following:

$$\rho^{(1)}{}_k = \int I^{(1)}(\lambda) R_k(\lambda) d\lambda = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(1)}{}_j g_{ijk} , \quad (3)$$

and

$$\rho^{(2)}{}_k = \int I^{(2)}(\lambda) R_k(\lambda) d\lambda = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(2)}{}_j g_{ijk} , \quad (4)$$

with $k = 1 .. p$.

For the mutual reflection color signal, we first substitute the basis function expansions yielding

$$I_m^{(2)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(2)}{}_j E_i(\lambda) S_j(\lambda) + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \varepsilon_i \sigma^{(1)}{}_j \sigma^{(2)}{}_{j'} E_i(\lambda) S_j(\lambda) S_{j'}(\lambda) ,$$

where we have abbreviated α_{12} simply as α. Now multiplying by $R_k(\lambda)$ and integrating, we find the final equation in our set of equations to solve,

$$\rho^m{}_k = \int I^{(2)}(\lambda) R_k(\lambda) d\lambda + \alpha \int I^{(1)}(\lambda) S^{(2)}(\lambda) R_k(\lambda) d\lambda$$

$$= \rho^{(2)}{}_k + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \varepsilon_i \sigma^{(1)}{}_j \sigma^{(2)}{}_{j'} h_{ijj'k} , \quad (5)$$

for $k = 1 .. p$, where

$$h_{ijf'k} = \int E_i(\lambda)\, S_j(\lambda)\, S_f(\lambda)\, R_k(\lambda)\, d\lambda$$

and we have abbreviated $\rho^{m\,(2)}{}_k$ as $\rho^{m}{}_k$.

From the above equation we can see that it is possible to restrict our attention to measurements of $\rho^{(1)}{}_k$, $\rho^{(2)}{}_k$, and $\rho^{m}{}_k = \rho^{m\,(2)}{}_k$ on *one* side of the edge, since from the similar equation for $\rho^{m\,(1)}{}_k$ we can see that if a solution is found for α_{12} then α_{21} can be determined via a least-square solution involving the $\rho^{m\,(1)}{}_k$, $\rho^{m\,(2)}{}_k$, $\rho^{(1)}{}_k$, and $\rho^{(2)}{}_k$. So that if a solution for ε_i, $\sigma^{(1)}{}_j$, $\sigma^{(2)}{}_j$, and α_{12} can be found via measurements on one side of an edge in the image then a solution for α_{21} is also available, given measurements $\rho^{m\,(1)}{}_k$ in the mutual reflection region on the other side.

Using measurements $\rho^{(1)}{}_k$, $\rho^{(2)}{}_k$, and $\rho^{m}{}_k$, we are faced with solving the set of equations (3), (4), and (5) for the ε_i, for the reflectance weights of both surfaces $\sigma^{(1)}{}_j$ and $\sigma^{(2)}{}_j$, and also for the geometrical factor α.

In order that a solution be possible, there must be at least as many equations as unknowns. Counting components in the set $\{\varepsilon_i, \sigma^{(1)}{}_j, \sigma^{(2)}{}_j\}$ we have $m + 2n$ unknowns, and α adds one more. Equations (3), (4), (5) each provide p equations, for a total of $3p$ equations. Hence a solution is possible only if the condition $3p \geq m + 2n + 1$ holds. Therefore we cannot use $p = m = n = 3$ unless we impose a further constraint. Since we must only expect solutions to yield illumination and surface reflectance spectra up to a multiplicative constant, because the ε_i and σ_j occur in products in the color signal, we must make one further requirement on one of the ε_i or σ_j; we choose to set $\varepsilon_1 = 1$. I.e., for surfaces we determine relative reflectance but not brightness.

Now our condition reads $p \geq (m + 2n)/3$ and we have a nonlinear set of equations for ε_i, $\sigma^{(1)}{}_j$, $\sigma^{(2)}{}_j$, and α that is at least sufficient for a solution.

3. Implementation and Results

We reduce the solution of the nonlinear set of equations (3), (4), (5) to the solution of *linear* ones by breaking up the solution into a multi-stage algorithm. Starting with some reasonable initialization for ε_i, we iteratively solve in turn for $\sigma^{(1)}{}_j$, then $\sigma^{(2)}{}_j$, and finally ε_i. The condition $\varepsilon_1 = 1$ at each step of the iteration also forces α to converge.

An important consideration is that the algorithm should obtain a unique solution. Uniqueness and speed of convergence can be established by converting the problem to be solved into the form of a fixed-point problem (see, e.g., [11]). Choosing a particular algorithm corresponds to adopting a particular fixed-point version of a problem. There are often several possible options, and not all of them converge. In [6], we investigated the convergence property of our algorithm by explicitly converting it to the fixed-point form to which it is equivalent. Then the well-known Contraction Mapping Theorem was used to test uniqueness. We showed in [6] that this Theorem provides a useful preprocessing step for screening pixels that may lead to non-converging solutions. We show there that uniqueness is guaranteed except in those cases where the finite dimensional model itself represents the surfaces very poorly. The algorithm is stable; so long as the mutual reflection effect is indeed present, the solution proceeds to find α. As $\alpha \to 0$ we expect the method to fail; this would correspond to glancing angles of interreflected rays.

In [6] we applied the algorithm to several cases involving synthesized color signals composed from naturally

occurring lights and reflectances, as well as noisy signals. To make up inputs to the algorithm, we synthesized color signals based on the two-zone, infinite-bounce set of equations (1). The role of α is critical in understanding how accurate such two-zone color signals are. In the next section we investigate the geometrical significance of the factor α.

4. Configuration Factors

One may ask whether or not there generally exist points on a surface where we can effectively say that no mutual reflection effect exists. We can answer this question by examining the structure of the parameters α_{12} and α_{21}. The integral of *all* radiance reaching *all* points on the other surface as a fraction of the radiance from the first surface into the hemisphere above it is termed the *form-factor* in illumination engineering. Here we are instead interested in how much light reflected from one surface is intercepted at a *particular* spot on the second surface. A catalog of such *configuration factors* has been assembled in [12], and in general their values can be calculated for any geometry and are independent of wavelength for ideal diffuse reflectors.

4.1. Two-zone approximation

The set of equations (1) effectively embodies a wavelength-dependent radiosity calculation [7, 12] where each emissive term is replaced by the color signal $S^{(i)}(\lambda)E(\lambda)$ coming from each surface, independent of interreflection; and one patch is distinguished by the property that all configuration factors involving it are zero. Thus our algorithm effectively analyzes a one-bounce, two-zone radiosity exchange, where radiosity is defined as the net radiation leaving a surface (comprised of both emissive and reflective terms in graphics).

For illustrative purposes, consider for simplicity a relatively long (effectively semi-infinite) planar edge (Fig. 2). Let the fraction of light from the small area ΔA_1 intercepted by the upper surface A_2 divided by the the total amount of light from ΔA_1 into the hemisphere above it be F_{12}. Calculation of the quantity F_{12} will yield as well F_{21}, the fraction of light from A_2 intercepted by ΔA_1, because of the reciprocity relation for configuration factors [12]

$$A_2 F_{21} = \Delta A_1 F_{12} .$$

The configuration factor for Fig. 1 is found to be [12]:

$$F_{21} = \frac{\Delta A_1}{2 A_2} [\, 1 + \frac{\cos\beta - X}{(X^2 + 1 - 2X\cos\beta)^{1/2}} \,] ,$$

with the dimensionless distance from the vertex defined as $X \equiv x/\xi$. We show in [6] that this corresponds to a two-zone α factor

$$\alpha_{21} = \frac{1}{2} [\, 1 + \frac{\cos\beta - X}{(X^2 + 1 - 2X\cos\beta)^{1/2}} \,] . \qquad (6)$$

The factor α_{21} is plotted in Fig. 3 for representative opening angles β. One can see that as the position X from the vertex increases, the factor α_{21} decreases monotonically with limit 0. Typical values with β=120° are 0.107 for X=0.6 and 0.028 for X=2.0. For flat surfaces, Fig. 3 shows that there are areas where α is relatively small provided one surface is sufficiently wider than the other. For real surfaces, particularly convex ones, there will certainly be many sites classifiable as being far from edges or occlusions.

The configuration factor in this example gives the contribution of light from surface A_2 reaching ΔA_1; however, it is important to note that α_{21} results from the assumption that A_2 has uniform color signal $I(\lambda)$. I.e., the whole side forms a single zone with the exception of one zero-mutual-reflectance point. Relaxing this assumption leads to a different factor, requiring a full radiosity calculation. We now consider this more general case.

4.2. Multi-zone method

The radiosity method can be used to generate more detailed variation in the color signal from each side of the edge. This method consists (see [7] and [12], p. 248) of subdividing each side into a number of sub-zones and summing up the contribution of each.

Here we generate a full, multi-zone analysis of a planar edge, using the radiosity method and then, in Section 5.2, test our approximate surface reflectance recovery algorithm on each pixel along the surface.

Consider the edge shown in Fig. 4. Here we vary x on A_1 from 0 to X and vary ξ on A_2 from 0 to Ξ. The radiosity method consists of a discretization of a set of integral equations formed by an infinite subdivision of each surface. In general, for a scene with surfaces $\mu = 1 \,..\, M$, if an element on surface μ is at $\vec{r}_\mu$ then the color signal from that spot, including mutual reflection, is

$$I_m^{(\mu)}(\vec{r}_\mu,\lambda) = S^{(\mu)}(\lambda)E(\lambda) + S^{(\mu)}(\lambda)\sum_{\nu=1}^{M}\int_{A_\nu} I_m^{(\nu)}(\vec{r}_\nu,\lambda)\, dF_{\mu\nu}(\vec{r}_\mu,\vec{r}_\nu) \tag{7}$$

where $\vec{r}_\nu$ parameterizes A_ν and $dF_{\mu\nu}(\vec{r}_\mu,\vec{r}_\nu)$ is a surface-element to surface-element configuration factor.

In the above, the order of subscripts in $dF_{\mu\nu}$ is important. It follows from the reciprocity relation for configuration factors ([12], p. 178) $dA_\nu\, dF_{\nu\mu} = dA_\mu\, dF_{\mu\nu}$ and from considering the incoming power $P^{(\mu)}_{in}$ arriving at dA_μ:

$$P^{(\mu)}_{in} = \sum_{\nu=1}^{M}\int_{A_\nu} dA_\nu I_m^{(\nu)}(\vec{r}_\nu)\, dF_{\nu\mu}$$

$$= \sum_{\nu=1}^{M}\int_{A_\mu} dA_\mu I_m^{(\nu)}(\vec{r}_\nu)\, dF_{\mu\nu} \,.$$

Dividing by dA_μ leads one to eqn.(7).

In Fig.4, an edge is formed by two flat semi-infinite strips of finite width, joined by an angle β. The configuration factor from one infinitesimal area to another in given by [12]

$$dF_{x\xi} = \frac{x\xi \sin^2\beta}{2\,(x^2+\xi^2-2x\xi\cos\beta)^{3/2}}\, d\xi = \phi(x,\xi)\, d\xi$$

where we denote the finite part of the above factor by $\phi(x,\xi)$.

Subdividing each strip into an infinite number of patches then leads to the set of integral equations

$$I_m^{(1)}(x,\lambda) = S^{(1)}(\lambda)\, E(\lambda) + S^{(1)}(\lambda) \int_{\xi=0}^{\Xi} I_m^{(2)}(\xi,\lambda)\, \phi(x,\xi)\, d\xi \tag{8}$$

$$I_m^{(2)}(\xi,\lambda) = S^{(2)}(\lambda)\, E(\lambda) + S^{(2)}(\lambda) \int_{x=0}^{X} I_m^{(1)}(x,\lambda)\, \phi(\xi,x)\, dx \tag{9}$$

Substituting (9) into (8) yields

$$I_m^{(1)}(x,\lambda) = S^{(1)}(\lambda)\, E(\lambda) + S^{(1)}(\lambda)\, S^{(2)}(\lambda)\, E(\lambda) \int_{\xi=0}^{\Xi} \phi(x,\xi)\, d\xi \tag{10}$$

$$+ S^{(1)}(\lambda)\, S^{(2)}(\lambda) \int_{\xi=0}^{\Xi} \int_{x'=0}^{X} I_m^{(1)}(x',\lambda) \phi(x',\xi)\, \phi(x,\xi)\, dx'\, d\xi$$

The first term represents the no-bounce color signal. The second term gives the one-bounce contribution from all of surface 2 to dx on surface 1.

We have

$$\int_{\xi=0}^{\Xi} \phi(x,\xi)\, d\xi = \frac{1}{2}\left[\, 1 + \frac{\Xi\cos\beta - x}{(x^2 + \Xi^2 - 2x\Xi\cos\beta)^{1/2}}\,\right] \equiv \zeta(x)\,,$$

similar to eqn.(6). We denote this integral by $\zeta(x)$. This is a finite surface to surface-element configuration factor, as given in tables for this geometry. The difference between α_{21} and ζ is that whereas ζ takes into account all the $I(\lambda)$ changes from one side to the other for each sub-zone, α_{21} effectively lumps these changes together by assuming each side consists of only one zone. The last term in eqn.(10) gives the element-to-element color bleeding and is usually not explicitly shown in works on interreflection. It retains the integral equation character of eqns.(8) and (9).

From eqn.(10) one sees that each wavelength band is blended differently-- it is not in general possible to calculate color bleeding effects independent of the specific illumination and surface reflectance spectra in use.

5. Simulations

We tested the one-bounce, two-zone algorithm (see [6] for details) on many synthetic color signals made up according to the infinite-bounce, two-zone set of equations (1). In these simulations, and in those here, color signals are synthesized from naturally occurring illumination spectra and surface spectral reflectance functions. By passing these synthetic color signals through bandpass filters representing film or other sensor sensitivities, we determine sensor values ρ_k, which then become the input values for our algorithm, along with the proportion α of the light from one surface reaching the second surface. The algorithm is then used to solve for an estimate of the surface spectral reflectance on each side of the edge, the illumination spectrum, and the mutual reflection factor α.

5.1. Two-zone simulations

For known α_{12} and α_{21}, we constructed [6] a two-zone mutual reflection input signal and tested the algorithm by solving the approximate model equations (3,4,5). For simplicity we assumed a geometry leading to $\alpha_{12} = \alpha_{21}$ to generate the input values of ρ^m from eqns. (1). For surfaces we chose pairs of reflectances from the collection of 370 surface spectral reflectance measurements catalogued by Krinov [13]. For illuminations we used the standard daylights tabulated by Judd et al. [10] as well as those measured by Dixon [14].

Applying the algorithm of Section 3 to these color signals, we calculated the percentage error in the fit of the whole spectrum to the actual reflectance spectrum. Comparing this error to the the error in the underlying finite-dimensional model, we found an average ratio of ~1.3 times the error in the underlying finite-dimensional model. For noisy signals, consisting of color signals derived from eqns.(1) with additive noise, we found that for an addition of 5% noise the ratio of errors rose to about 2.

These results show that the one-bounce, two-zone approximation of mutual reflection works reasonably well when presented with synthetic color signals made up according to an infinite-bounce, two-zone model of interreflection. One important point in the algorithm is the necessity of being able to find points where the mutual reflection drops off to essentially zero. Another point is how well the algorithm performs on more realistic edges. To test these two points we consider another edge, still synthetic, but produced with a complete, multi-zone model for interreflection.

5.2. Multi-zone simulation

Using the multi-zone model developed from the radiosity method in Section 4.2, we consider again the planar edge of Fig. 4. Taking as example spectra, samples #54 and #248 from Krinov's catalogue [13] of naturally occurring reflectances, using incident illumination D65, and choosing an opening angle β=135°, we arrive at the color-blended edge shown in Fig. 5. Here we show the variation in RGB on surface 2 as ξ is increased away from the edge. For comparison we show also the RGB resulting from $S^{(2)}(\lambda)\, E(\lambda)$ with no interreflection. To show that color, as well as intensity, is changing along the surface, we plot the ratios R/B and G/B in Fig. 6.

Now carrying out our algorithm with the color-blended edge as input, we find that reflectances on each side are recovered with a reasonable degree of accuracy. Except at pixels very close to the edge, the algorithm obtains a reflectance $S^{(2)}(\lambda)$ that is virtually independent of position. The error with respect to the actual reflectance spectrum is 4.37%; this translates into a color difference ΔE of 3.54 units in CIELUV 3-dimensional uniform color space (see [15]). For surface 2, the best least squares approximation using the finite dimensional model with dimensionality 3 gives an error of ΔE=3.25; the mutual reflection method is seen to contribute only a small additional error. For illumination, the error of the recovered spectrum is virtually zero.

The recovered values of α vary as shown in Fig. 7. Recall that the algorithm finds a one-bounce α that accounts for all the increased intensity resulting from multi-bounce interreflection in one lumped parameter. Thus one could expect α to be larger than ζ. In general, α will vary with the intensity in order to compensate for holding $\varepsilon_1 = 1$. However, the two factors α and ζ may have different units, since the units of α are dependent on the normalization chosen for the illumination basis functions $E_i(\lambda)$, as a result of the condition $\varepsilon_1 = 1$ in the algorithm. As well, since the two-zone model assumes α is identically zero at x=X and ξ=Ξ, the α recovered by the algorithm should drop to

zero at these locations.

Therefore, one cannot generally compare α with the factor ζ in an absolute sense. However, here we have a priori knowledge of the illumination intensity and can scale α to have the same units as ζ. In Fig. 7 we show α scaled by multiplication by a factor $\varepsilon_{1,exact}$, where $\vec{\varepsilon}_{exact}$ is the best least squares fit to the input illumination. Also, since the value of $\rho^{(2)}$ that is measured is not a no-bounce RGB but actually contains some interreflection contribution as in eqn. (10), the factor α should really be compared to a factor $\zeta(\xi) - \zeta(\Xi)$. We can see from Fig. 7 that α closely approximates this function but is slightly larger because it takes into account multiple reflections. As it must, α goes to zero at the point where $\rho^{(2)}$ is measured. In general a determination of α would give a curve similar to Fig. 7, but scaling and offset information would not be available.

6. Conclusions

In conclusion, we can evaluate the algorithm as being reasonably robust. The method has the potential for eliminating the mutual reflection effect and in most cases recovers a good approximation of complete spectral information. Sensor values from only three pixels need be measured to obtain a result and the method will work on surfaces that do not meet the surface complexity conditions needed to make Maloney and Wandell's algorithm work. The elimination of mutual reflection is a desirable feature in a vision system since it is an effect which introduces extra complexities; e.g., it is known to create problems in shape from shading schemes [4, 5].

The algorithm still works even if both surfaces have the *same* color, with the same or differing brightness. This case might correspond to a folded piece of uniform material. The recovery of α here gives shape information near the fold. Why the method works in this one-color case is that the reflected light is a different color from the illumination, and the extra measurements of sensor values supply enough information to disentangle the two.

The results presented here show that, in situations in which mutual reflection is present, measurements of the sensor values taken separately from the two different interreflecting surfaces, as well as from within the mutual reflection zone, convey sufficient information to allow complete recovery of the two surface spectral reflectance functions as well as of ambient illumination. We also recover the factor α, which summarizes information about the relative positions of the two surfaces. The recovery algorithm imposes the simplified one-bounce, two-zone model of mutual reflection. We can conclude that this simplified model is sufficient for the recovery algorithm since the results in Fig. 4 show that α drops off sufficiently fast in a real, multi-zone edge to guarantee finding locations with negligible mutual reflection.

For situations in which mutual reflection is everywhere present, one could abandon the insistence on being able to find α=0 sites by encompassing all available pixel information in a variational calculus scheme minimizing deviations from the one-bounce model. Work is being pursued in this direction.

References

1. L.T. Maloney, *Computational Approaches To Color Constancy*, PhD dissertation, Stanford University, 1985, Applied Psychology Laboratory

2. L.T. Maloney and B.A. Wandell, "Color constancy: a method for recovering surface spectral reflectance", *J. Opt. Soc. Am. A*, Vol. 3 1986, pp. 29-33.

3. B.A. Wandell, "The synthesis and analysis of color images", *IEEE Trans. Patt. Anal. and Mach. Intell.*, Vol. PAMI-9 1987, pp. 2-13.

4. D. Forsyth and A. Zisserman, "Mutual illumination", *Proceedings: Computer Vision and Pattern Recognition*, IEEE, 1989, To be published

5. D. Forsyth and A. Zisserman, "Shape from shading in the light of mutual illumination". Preprint, Robotics Research Group, Department of Eng. Sci., Oxford University

6. B.V. Funt, M.S. Drew, and J. Ho, " Color constancy from mutual reflection", Tech. report CSS/LCCR TR 89-02, Simon Fraser University School of Computing Science,1989.

7. C.M. Goral, K.E. Torrance, D.P. Greenberg, and B. Battaile, "Modeling the interaction of light between diffuse surfaces", *Computer Graphics*, Vol. 18 1984, pp. 213-222.

8. J. Cohen, "Dependency of The Spectral Reflectance Curves of The Munsell Color Chips", *Psychon. Sci.*, Vol. 1 1964, pp. 369-370.

9. J.P.S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic spectra of Munsell colors", *J. Opt. Soc. Am. A*, Vol. 6 1989, pp. 318-322.

10. D.B. Judd, D.L. MacAdam, and G. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature", *J. Opt. Soc. Am.*, Vol. 54 August 1964, pp. 1031-1040.

11. R.L. Burden, J.D. Faires, and A.C. Reynolds, *Numerical Analysis*, Prindle, Weber & Schmidt, 1981.

12. R. Siegel and J.R. Howell, *Thermal Radiation Heat Transfer*, Hemisphere Pub. Corp., 1981.

13. E.L. Krinov, "Spectral Reflectance Properties of Natural Formations", *Technical Translation TT-439, National Research Council of Canada*, 1947.

14. E.R. Dixon, "Spectral distribution of Australian daylight", *J. Opt. Soc. Am*, Vol. 68 1978, pp. 437-450.

15. Wyszecki, G. and Stiles, W. S., *Color Science: Concepts and Methods, Quantitative Data and Formulas*, Wiley, New York, 1982.

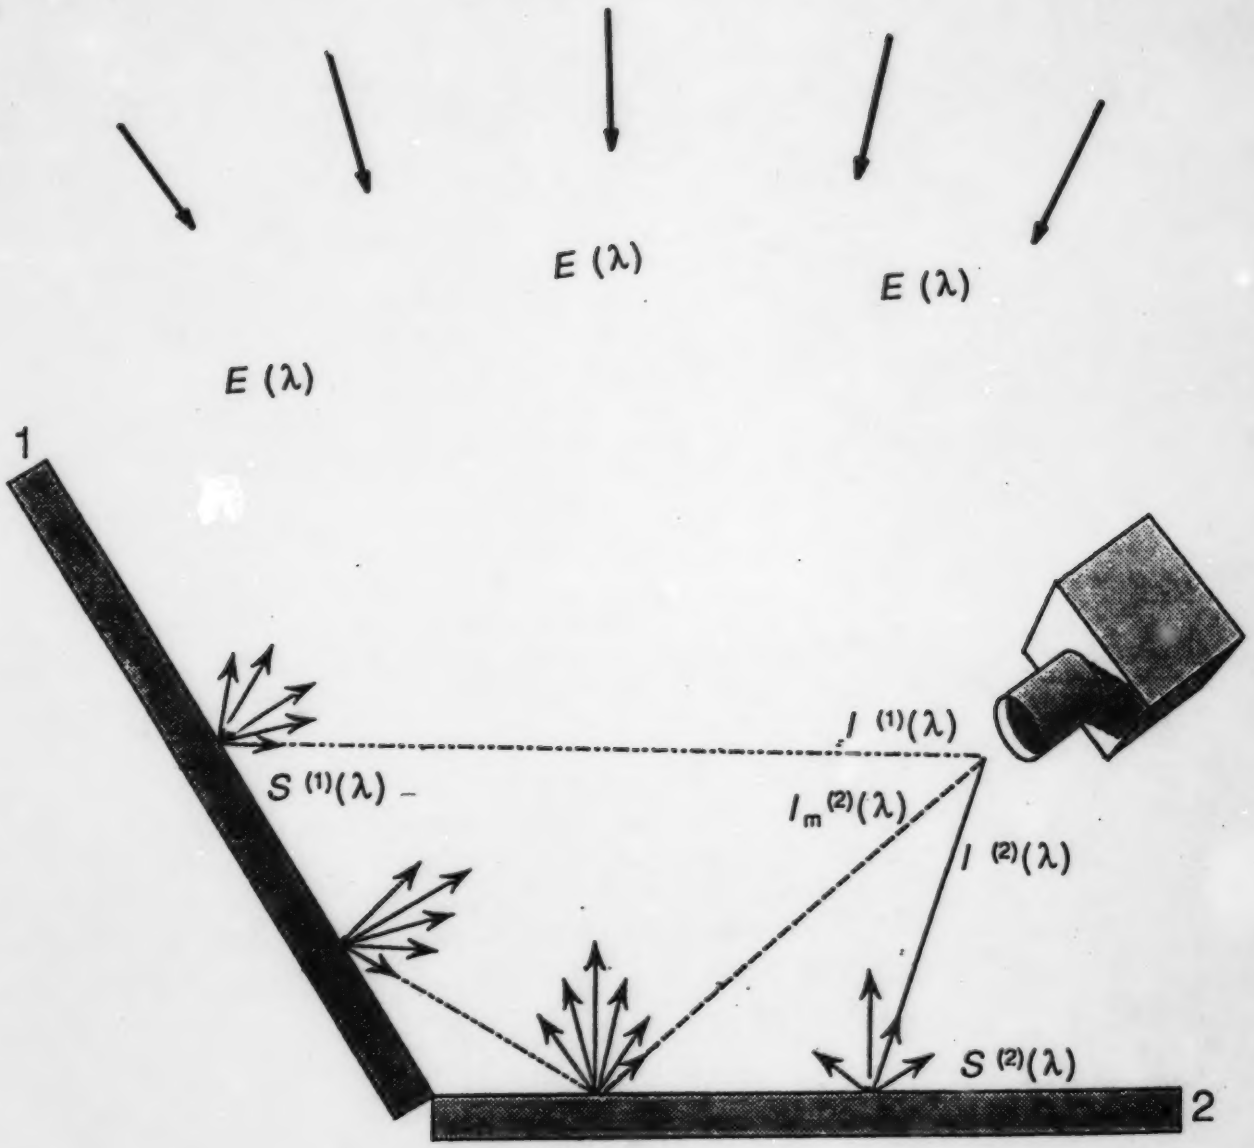

Fig.1. Illumination with spectral power distribution $E(\lambda)$ impinges on surfaces 1 and 2, which have surface reflectance functions $S^{(1)}(\lambda)$ and $S^{(2)}(\lambda)$ respectively. Far from the edge surface 1 reflects color signal $I^{(1)}(\lambda)$ towards the camera and similarly surface 2 reflects $I^{(2)}(\lambda)$. Close to the edge $E(\lambda)$ is augmented by light reflected from the other surface. The color signals reflected from the surfaces near the edge are $I_m^{(1)}(\lambda)$ and $I_m^{(2)}(\lambda)$.

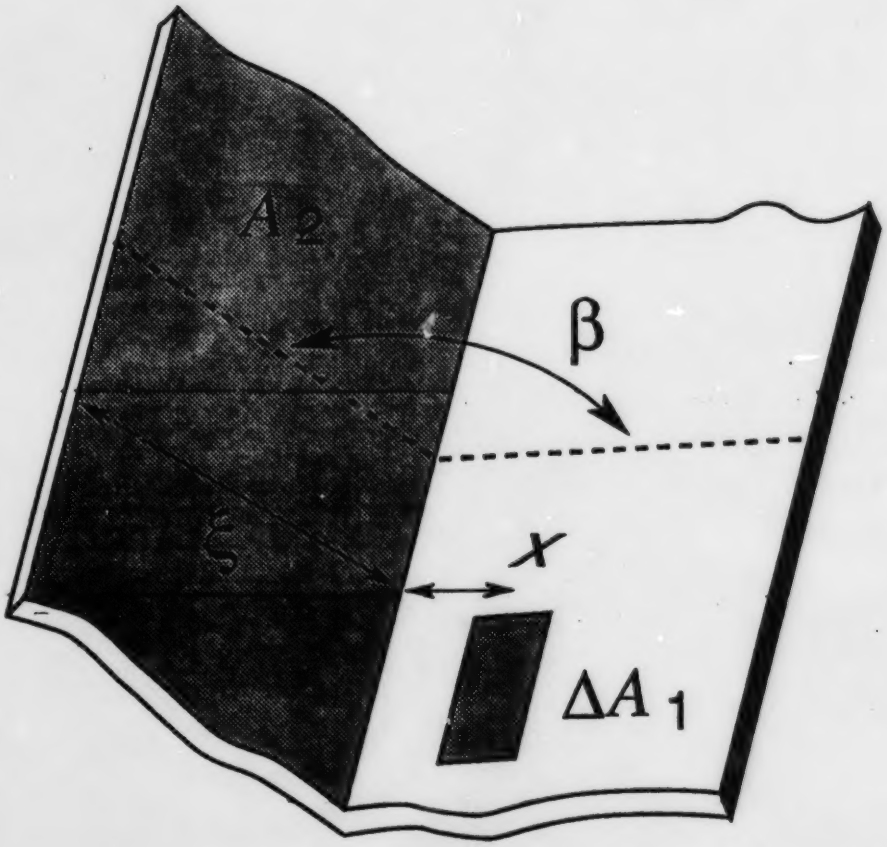

Fig.2. A long edge formed from flat surfaces. The opening angle is β. ξ is the width of surface 2, with area A_2. The small surface area ΔA_1 is at distance x from the edge.

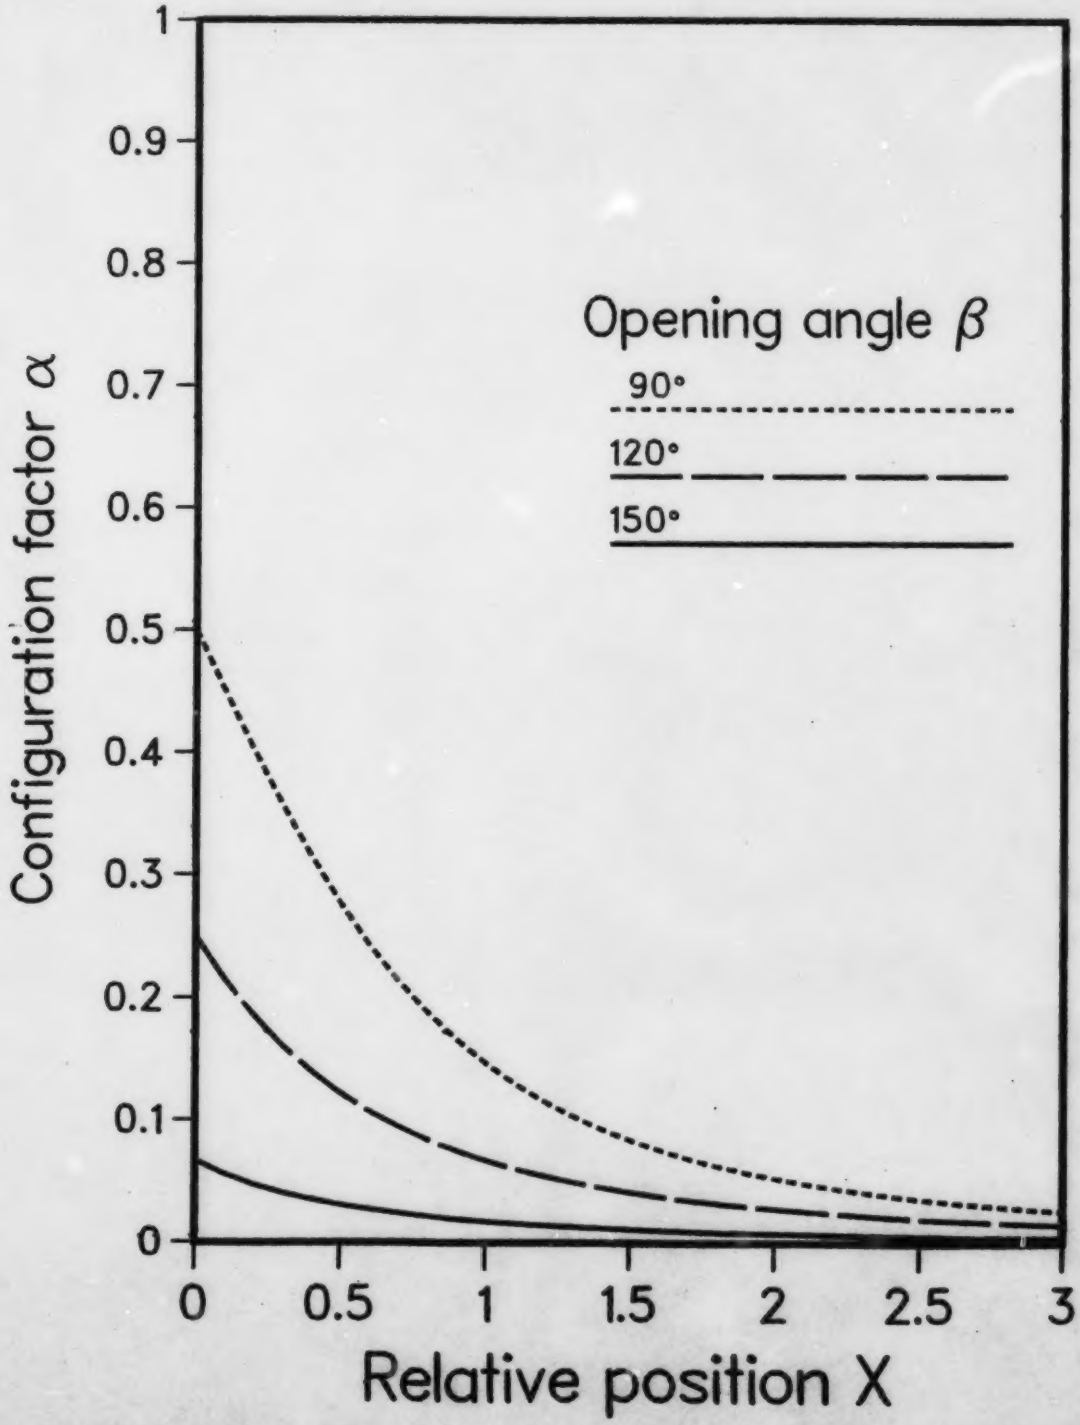

Fig.3. Configuration factor α as a function of $X = x/\xi$ for several opening angles β.

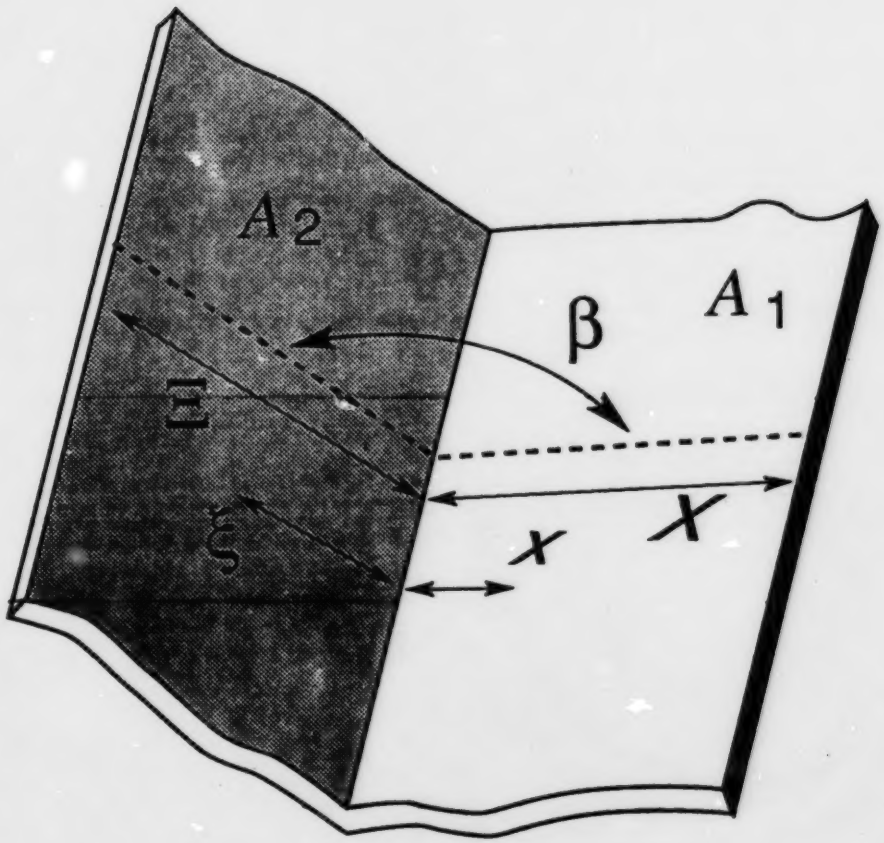

Fig.4. Geometry of a plane edge for the multi-zone model.

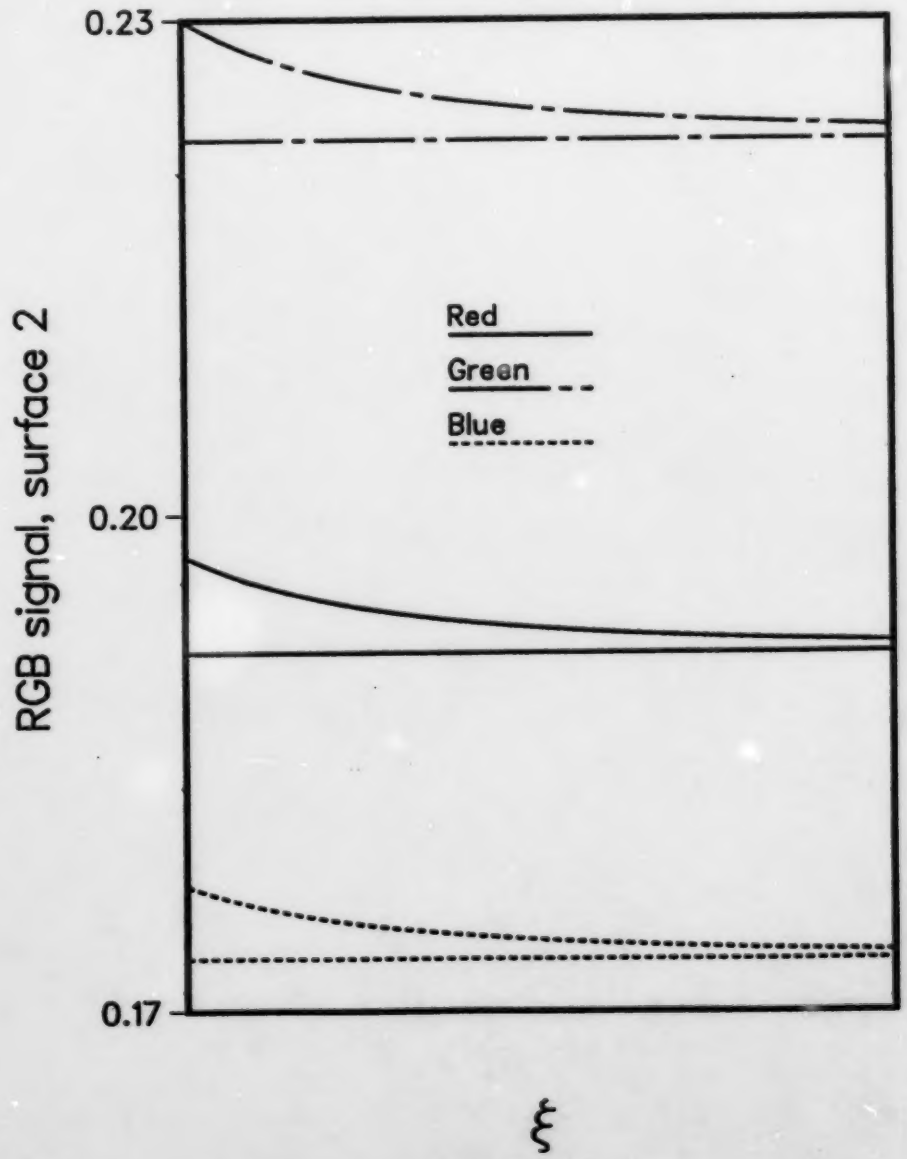

Fig.5. RGB variation on one side of an edge due to mutual reflection.

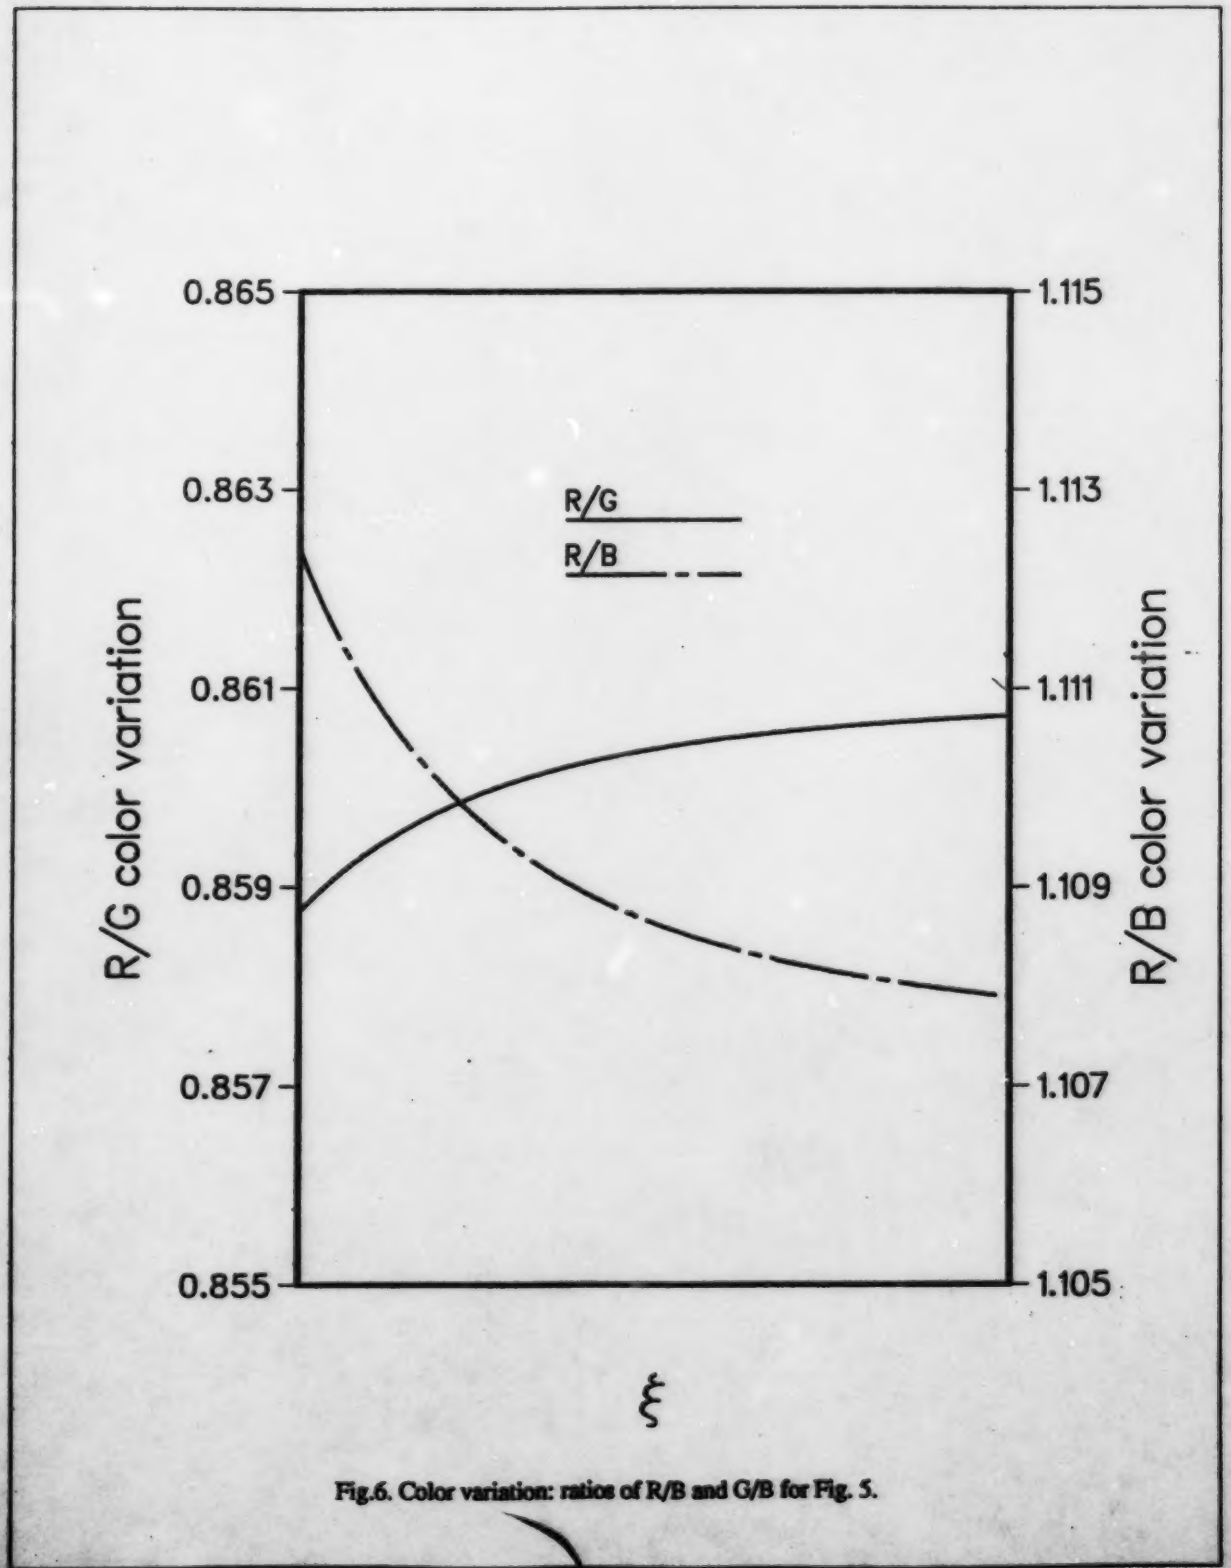

Fig.6. Color variation: ratios of R/B and G/B for Fig. 5.

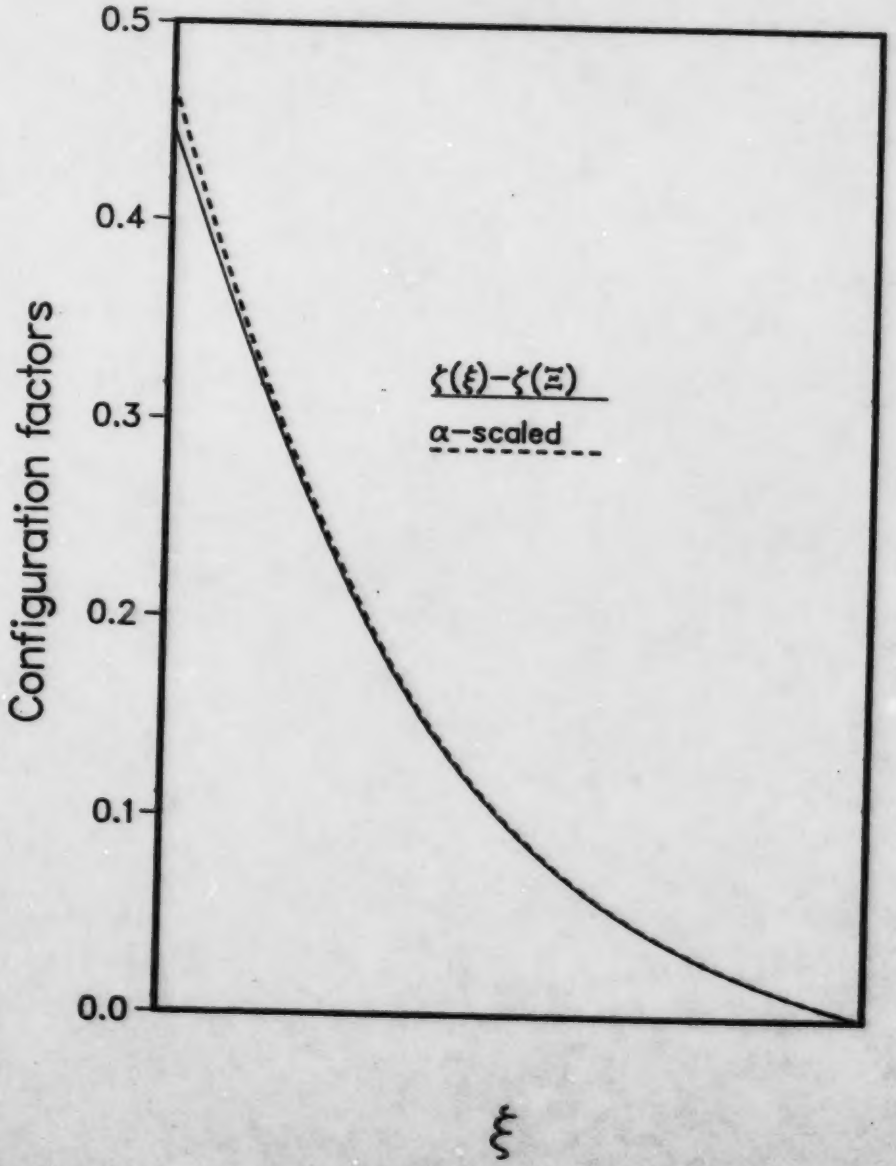

Fig.7. Variation of one-bounce, two-zone α factor, where the sensor RGB inputs are derived from a full multi-zone model. One-bounce, uniform color factor ζ is shown for comparison.

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

• CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.